

# Exhaustive Geographic Search with Mobile Robots Along Space-Filling Curves<sup>1</sup>

Shannon V. Spires  
Steven Y. Goldsmith

Advanced Information Systems Laboratory  
Sandia National Laboratories  
Albuquerque, New Mexico USA  
*svspire@sandia.gov, sygolds@sandia.gov*

**Abstract.** Swarms of mobile robots can be tasked with searching a geographic region for targets of interest, such as buried land mines. We assume that the individual robots are equipped with sensors tuned to the targets of interest, that these sensors have limited range, and that the robots can communicate with one another to enable cooperation. How can a swarm of cooperating sensate robots efficiently search a given geographic region for targets in the absence of a priori information about the targets' locations? Many of the "obvious" approaches are inefficient or lack robustness. One efficient approach is to have the robots traverse a space-filling curve. For many geographic search applications, this method is energy-frugal, highly robust, and provides guaranteed coverage in a finite time that decreases as the reciprocal of the number of robots sharing the search task. Furthermore, control is inherently decentralized and requires very little robot-to-robot communication for the robots to organize their movements. This report presents some preliminary results from applying the Hilbert space-filling curve to geographic search by mobile robots.

## Introduction

The idea of using swarms of cooperating robots to solve various physical problems has received much attention recently [CFKM95], and has recently become cost-effective and practical with the advent of less expensive hardware and with new analysis techniques from complexity theory, chaos theory, and nonlinear dynamics. One problem that is particularly applicable to robot swarms is exhaustive geographic search. Exhaustive geographic search asks that we develop a complete map of all phenomena of interest within a defined geographic area, subject to the usual engineering constraints of efficiency, robustness, and accuracy [GR98b], [CP97], [HTL96], [KOAY95].

One example of such a search problem is that of finding buried land mines. It is possible to build a robot that has the requisite sensors and navigation apparatus to do this job, thus removing humans from an extremely high-risk activity. But a single robot is likely to be expensive and subject to damage from the mines, rendering it useless. If less expensive robots could be built in quantity, it would be better to build a swarm of such robots that could cooperate to locate all the mines within a given geographic region. With the right programming, many robots could do the job faster than one (ideally in  $1/m$  the time, if  $m$  robots participate). The search process would be more robust because if a robot is damaged by a mine, others could take over its work.

This report presents an efficient, robust, cooperative search algorithm that allows multiple robots to be applied to such a mission. The primary insight is that exhaustive coverage of a one-dimensional path with multiple robots is substantially easier and more robust than coverage in two (or more) dimensions. This algorithm transforms a two-dimensional search task into one-dimensional search, with the result that two dimensions can be effectively searched with the attendant advantages of one-dimensional coverage. The transformation is accomplished by having the robots traverse a minimum-length space-filling curve, dividing the curve among  $m$  robots.

## Problem Constraints

The nature of the problem is such that there is no *a priori* knowledge of the locations of the targets (the targets are the land mines in the example above); we merely know that there *may* be one or more targets within a given geographic region. Our mission is to discover definitely their total number, their locations, and optionally to further characterize them along other dimensions. The requirements and constraints of the task are as follows:

- Each robot has a sensor apparatus which is adequate for detecting the targets of interest, but its range is limited to an area much smaller than the overall region of interest. We assume that all the robots have the same sensor range, or that they can all be made aware of the minimum sensor range of any of the robots.
- Each robot can reliably communicate with the other robots with some probability  $P_c$ . The search mission should

<sup>1</sup> Sandia is a multiprogram laboratory operated by Sandia Corporation, a Lockheed Martin Company, for the United States Department of Energy under Contract DE-AC04-94AL85000.

be completable even as  $P_C$  approaches 0, although it may take longer.

- Each robot has a finite energy supply (e.g. batteries or a fuel cell).
- The robots “know” when their job is finished.
- The robots do not collide.
- The robots stay within the bounds of the search region.
- A robot’s individual search area should not overlap<sup>2</sup> that of another robot, but the entire search area must be covered (hence the use of the term *exhaustive* search).
- The search should still be completable even if one or more robots is damaged or otherwise fails during the mission. As long as at least one robot remains alive, the search should complete, although it may take longer.
- We must accomplish the search within a finite time.
- The search time should decrease as the number of robots participating increases.
- Control of the robots is completely decentralized; there is no central coordinating entity. The robots must be able to behave autonomously. Isolated robots must still be able to perform useful work.

## Geographic Search

In general, a geographic search will take place in three phases: (1) Initial configuration; (2) Search; (3) Terminal configuration [GR98a]. First, the robots must organize themselves into a coherent, communicating group within the region. They divide the area into subregions and assign each robot one or more subregions. Next, the search task proper is conducted. Finally, the robots determine through consensus that the region has been searched. At this point, the robots take some terminal action.


To maximize efficiency, two considerations are immediately obvious:

A) *Minimal configuration energy*. We don’t want the robots to expend a great deal of energy in the initial configuration phase because no useful work is accomplished until the search phase. During initial configuration, the robots must expend energy communicating with each other and moving into their initial positions. These energy expenditures must be minimized.

B) *Optimal search coverage*. During the search phase, we must ensure that the entire search area is covered in a finite time, preferably proportional to the reciprocal of the number of robots involved.

To optimize (A), we could have the robots perform no initial configuration at all, with each simply searching the area in its immediate vicinity. No initial communication and no initial movement take place to coordinate with the other robots. We assume the robots are initially randomly distributed on the search region (because they were dropped out of an airplane, for example). Each robot could perform a random walk [W97a], spiral outward, or use some other autonomous algorithm. This defeats the purpose of multi-robot collaboration. Since there is no coordination, areas already searched can be retraced. Spiraling outward would guarantee that the search area was covered in a finite time, but using more than one robot would not necessarily decrease that time. A random walk is even worse—it wouldn’t even guarantee coverage in a finite time.

Optimizing (B) means having the robots subdivide the search area equally and each search its agreed-upon subregion. A good deal of initial collaboration and movement would generally be needed here, especially if the robots start in random locations and must move to an organized initial configuration. An example of such an initial configuration is shown in Figure 1, where eight robots (triangles) have lined up in a column along the left side of the search region and plan to march to the right. Each search subregion is a horizontal strip.



**Figure 1:** Robots (triangles) aligned in preparation to march to the right

This accomplishes a maximally-efficient search, but at the cost of a less efficient initial configuration. In addition, this search assumes all the robots are highly reliable; it’s not very robust if one or more robots dies during the search. (We’ll discuss why later in the report.) The dilemma is thus to find the optimal tradeoff between the initial configuration phase and the search phase that makes both acceptably efficient, while ensuring the robustness of the

<sup>2</sup> There may be cases where, for reasons of higher sensor reliability, we explicitly *do* want them to overlap.

mission.


We believe that searching along a space-filling curve solves the dilemma nicely: it requires very little initial configuration while also providing an efficient, robust search. Other work in coverage [CP97, HTL96] addresses obstacle avoidance (currently a weak point of our algorithm) but does not address efficient coverage with multiple robots or robustness. [KOAY95] discusses efficient multiple-robot sweeping but control is not distributed and robustness is not addressed.

We will digress a bit here to describe space-filling curves in general.

### Space-Filling Curves


A space-filling curve [S94] is a one-dimensional curve which passes through every point of a given N-dimensional region. It is thus a one-to-one mapping between an N-dimensional and a 1-dimensional space. Such curves can be constructed to fill regions of any dimensionality, but in our case we're mainly interested in two-dimensional regions, corresponding to a geographic area on the surface of the earth. Two-dimensional regions are also much easier to illustrate, and for those reasons the remainder of this report will concern only curves that fill 2-d regions. The reader should bear in mind, however, that the constructions and algorithms described herein can easily be extended to regions of more dimensions.<sup>3</sup>

Of course, a true space-filling curve is an ideal mathematical entity and can only be approximated in the real world. If we divide our region of interest into subregions of finite size, rather than points, we can easily draw a curve which passes through each subregion. In Figure 2, for example, a square region is divided into four subregions. The dark inverted U-shaped curve passes through the center of each subregion.



**Fig. 2.** First Order Hilbert curve


Subdividing further, into 16 regions (4 x 4) the curve becomes as in Figure 3.



**Fig. 3.** Order 2 Hilbert curve

Subdividing still further, into an 8 x 8 grid, Figure 4 shows a curve which passes through each of the 64 subregions:

<sup>3</sup> A 3-d region [G97], [SLP83] might be appropriate for search by a swarm of undersea [HTL96] or flying robots, for example.



**Fig. 4.** Order 3 Hilbert curve (gray arc at bottom is explained in the text)

The family of curves illustrated in Figs 2-4 is called the *Hilbert* [W97b] curve after David Hilbert, their discoverer. There are many families of space-filling curves, but the Hilbert has some especially nice properties. For example, two points that are “close” in 2-d space are more likely to be “close” in Hilbert space than with other space-filling curves [MJFS96].

Notice how every deeper subdivision of the curve contains four copies of the entire previous curve, suitably rotated and reflected, with some straight segments added to ensure continuity. The curve is thus geometrically self-similar, as space filling curves frequently are. The *order* of the curve  $D$  determines the number of subregions it passes through and consequently, the total length of the curve:

$$\# \text{ subregions} = (2^D)^2 = 2^{2D} = \text{Length}_{\text{curve}} + 1$$

Thus Figure 2 is a first order curve, or  $D=1$ , because it occupies 4 subregions and is 3 units long. Figure 3 has  $D=2$ , and Figure 4 has  $D=3$ . To fill every point in a region, a curve would have to be  $D=\infty$ .

The fact that the curve is a geometrically self-similar fractal implies that there is a recursive algorithm for its computation. This is indeed the case [McW97], [PLF91]. However, there is also a non-recursive algorithm for computing a given Hilbert curve which is more useful in our robotic swarm application.

If we wanted to identify a particular subregion of a Hilbert-traversed space—for example, the shaded subregion in Figure 4—we could do it in two different ways. We could specify its  $[x,y]$  location as  $[3,5]$  (where  $[0,0]$  specifies the lower left corner subregion marked with the letter B in the figure), or we could specify its position along the Hilbert curve as the scalar value 28, which indicates that it is the 28th subregion encountered along the curve measuring from the beginning of the curve at subregion B. Subregion B is location 0 in this frame of reference. There is a straightforward algorithm to convert the 2-space coordinates of real space into a 1-space Hilbert coordinate and an equally simple algorithm to convert from 1-space back into 2-space [B69]. Even though the dimensionality changes during this conversion, no information is lost because the total number of possible bits in  $x$  or  $y$  is  $D$ , while the total number of bits needed to specify the 1-space Hilbert coordinate is  $2D$ . The conversion algorithms merely interpret the coordinate bits in different ways.

### Search Along a Space-Filling Curve

By using 1-space to 2-space conversion algorithms, it is easy to program a robot to follow a space-filling curve. Assuming the robot is able to deduce its current  $x,y$  location (via Global Positioning Satellite, inertial sensors, etc.), the algorithm is shown in Table 1:

$x_0, y_0 \leftarrow \text{current location}$
$H \leftarrow \text{Convert-xy-to-Hilbert}(x_0, y_0)$
$H \leftarrow H + 1$
$x_1, y_1 \leftarrow \text{Convert-Hilbert-to-xy}(H)$

**Table 1:** Hilbert traversal algorithm.  $x_1, y_1$  represents the robot’s new location.

In practice, we perform the operation  $H \leftarrow H + 1 \text{ modulo } 2^{2D}$ , which causes the robot to return to the beginning point of the curve after it reaches the end; this closes the curve and makes it a topological circle instead of a line segment. The gray curved arc at the bottom of Figure 4 reflects this closing.

If we now give the robot a sensing mechanism whose sensory range is at least as large as a single subregion, by traversing the curve it can now search the overall region in time proportional to the number of subregions,  $2^{2D}$ .

Notice that this algorithm works regardless of where on the curve the robot starts; it need not start at the  $H=0$  point of the curve. This is one reason why the recursive algorithm is not used in the robotic application; it would be difficult to initialize the recursive control stack properly to allow traversal to begin at an arbitrary point. Starting at an arbitrary point becomes valuable when more than one robot is involved. If we have two robots, we could start one at the beginning of the curve (point B in Figure 4) and the other at the halfway point of the curve (point H in Figure 4); the resulting traversal of the curve takes place in half the time it would take with one robot. In general, if we start with  $m$  robots, we can search the space in time proportional to  $2^{2D}/m$ , provided we initially space the robots equally along the curve.

### Efficiency

We'll now examine how search along a space-filling curve meets efficiency criterion (A). If we assume an initial random configuration of the robots within the region, we can actually get away with almost no initial configuration movement whatsoever. We also assume that the search region is square<sup>4</sup> and that its bounds and the robots' individual sensor ranges are known by each robot *a priori*.

Configuration proceeds as follows:

- 1) Each robot senses its own initial location, either absolutely or relative to the search region. This is a fundamental assumption of the algorithm and should be easy to realize with Global Positioning Satellite receivers or some other mechanism.
- 2) Each robot independently computes the subregion size based on its sensor range.<sup>5</sup> Since each also knows the bounds of the search region, each also computes its own location relative to the search region (unless its location sensor already provided the relative location) and computes the order of the Hilbert curve based on the size of the subregions relative to the overall search region. Again, each robot will compute the same value for the order of the curve.
- 3) Each robot decides which subregion it is currently within by quantizing its  $x,y$  location relative to the overall region. It then moves to the center of that subregion. This is the only time its wheels need to turn during the initial configuration phase, and the maximum distance any robot will have to move is

$$\frac{s}{\sqrt{2}}$$

where  $s$  is the size of a subregion square measured along a side.

- 4) The robots broadcast their individual initial locations to each other so that each knows the starting locations of all the others. (How the robots reliably communicate this information is itself an interesting question but is beyond the scope of this report [G93]. The clustering properties of the Hilbert curve may be useful here in case the robots are too far apart to achieve total communication.)

Configuration is now complete and the actual search begins. Each robot proceeds to each subregion in turn along the Hilbert curve, using the algorithm in Table 1. At each subregion, the robot senses any targets of interest and checks to see if its next subregion is the starting subregion of any other robot. If so, it stops and broadcasts "I'm Done!" along with information about any targets it may have found along the way. When all the robots have reported in, the collective can by consensus decide to report back to a higher authority, perform other tasks, or move on as a group to search another region. If after some reasonable time one or more robots have not reported in, the robot immediately behind the nonreporting robot can autonomously decide to proceed into the nonreporting robot's search area, search it, and report "I'm Done!". Note that each robot can precompute its own expected time to complete its portion of the search, since it knows at the outset how much of the curve it must traverse. These timeout values can be communicated to allow the collective to detect such fail-stop [K95] conditions.<sup>6</sup>

This leads us into an analysis of the expected search time. Certainly the time required to achieve guaranteed exhaustive coverage of the search region is bounded. At worst, the search will take time proportional to the number of subregions in the search region (or equivalently, the length of the space-filling curve). The precise upper bound time,  $T_{UB}$ , can be calculated by multiplying the expected movement speed of a robot by  $s$  (the size of a subregion) by the total number of subregions. At best, the lower bound time will be  $T_{LB} = T_{UB}/m$ , with  $m$  the number of robots. The expected exhaustive coverage time  $T_E$  will equal  $T_{LB}$  when the robots are initially configured to be spaced equidistant from each other along the space-filling curve. But we're not doing that; we're just moving each to the center of the subregion it initially occupies.<sup>7</sup> In this kind of initial configuration,  $T_E > T_{LB}$  but is still


<sup>4</sup> Non-square regions and even general polygonal regions can be handled by extensions of the techniques presented here.

<sup>5</sup> We assume for now that all the robots have the same sensor range, and thus compute the same subregion size.

<sup>6</sup> The issue of Byzantine failures [K95], [LSP82], in which a robot pretends to cooperate but actually is a saboteur, must also be addressed in environments where malevolent robots may be present.

<sup>7</sup> We assume for the sake of simplicity that no two robots initially occupy the same subregion. For a real implementation, a resolution protocol would be needed to deal with this possibility.

nowhere near  $T_{UB}$ . Empirical results based on Monte Carlo simulation are shown in Figure 5. The graph clearly shows a reciprocal relationship between  $T_E$  and number of robots. It shows that with 20 robots, the expected time to exhaust the space is less than 20% of the time required with 1 robot. This is obviously not as good as the theoretical best case of  $T_{LB} = 5\%$ , but it's not bad for almost no initial configuration movement. We believe that with *slightly* more initial configuration movement,  $T_E$  can come much closer to  $T_{LB}$ . The nature of the particular space-filling curve being used becomes very important here. This is an open issue we are investigating.




**Fig. 5.**  $T_E$  (upper curve) as a percent of  $T_{UB}$  vs.  $m$ . ( $T_{LB}$  is shown for comparison purposes.) Random initial positions and minimal initial configuration movement. Empirical results.

Another point is that with the Hilbert curve, the search region must be divided into an integer power of 2 subregions (measured along a side). This can require the subregions to be significantly smaller than the sensor radius, thus increasing the overall distance traveled beyond that which would be required in the ideal case. Other curves which are not recursively subdivided by 2 but by other factors could remove this inefficiency, as could multi-level decomposition of the search region into macro- and micro- regions, where each type of region has a different coverage algorithm.

### Robustness

In the initial configuration shown in Figure 1, where a space-filling curve is not being used, what happens if one of the robots is defective, or dies during the march? Its strip of territory will not be searched. The other robots will have to reliably detect this and reorganize themselves to search the defective robot's area. But if the robots are searching along a space-filling curve, all the robots are topologically searching in one dimension along a single circle (Figure 6).



**Fig. 6.** Topological interpretation of search along a space-filling curve. Robots (triangles) traverse the curve clockwise.

Here, each robot eventually would search the entire space, because they are all searching along the same path. Ideally, each robot would stop when it encounters the starting point of the next robot along the path, with the result being that the space is searched in  $1/m$  the time it would normally take with 1 robot. But if one robot breaks down, the next robot behind can simply continue along its path without stopping, and the overall space still gets searched,

albeit in slightly longer time. No reconfiguration of the robots is necessary to accomplish this. Indeed, if as many as  $m-1$  of the robots all break, the search is still guaranteed to complete (assuming the surviving robot has enough energy reserves to search the entire space).

The search is also robust with respect to communication breakdowns. If all the robots can communicate with each other, the only groupwide communications needed are for initial configuration, and when a robot encounters the starting point of the next robot along the path<sup>8</sup>, it must announce to the collective "I'm Done!". When all robots have reported in, the collective knows the region has been completely searched. If one or more robots have not reported in by some time limit, the collective can deduce that they are dead and the next robot behind each dead robot can unilaterally decide (with no need for further communication with the group) to search the dead robot's area. If all communications fail (because of RF jamming, for example) the worst case is that the search will take as long as it would with only a single robot. Each robot will end up searching the entire space, and each will know to stop when it reaches *its own* starting point. The search will take a long time, but it's still guaranteed to complete. Thus communications failures can cause the search mission to take longer, but they cannot prevent it from being accomplished.

### Collision Avoidance

In any collaborative robotics system, collision detection and avoidance will be needed. Space-filling curves may make collision avoidance easier. When several robots follow a single space-filling curve, there is theoretically no chance of two robots' paths crossing, and therefore no need for dynamic collision avoidance. Realistically, two robots could collide if navigation errors are allowed to accumulate, if two adjacent robots along the curve travel at different speeds, or if the one in front dies. But it's possible that collision detection could be simplified because of the automatic separation inherent in the algorithm.

### Further Research

Much work still needs to be done exploring the general problem of robot traversal of space-filling curves. Some avenues of exploration include obstacle avoidance, formal expected exhaustion time analysis, handling robots with differing sensor ranges, non-square and non-rectangular search regions, and automatic determination of the search region based on energy stores and initial positions of the extrema robots in a cluster. Obstacle avoidance might be done by adapting the techniques described in [CP97], [HTL96], and [KOAY95]. Space-filling curves other than the Hilbert need to be investigated as well. One disadvantage of the Hilbert curve is the long straight path that connects its endpoints when it is closed. Other curves have the property that their beginning and end points are adjacent. Also, in cases where there is a penalty for turns, other space-filling curves may do the job with fewer turns than the traditional Hilbert. The Hilbert II curve [W97b] might be a candidate here.

And of course, the tradeoff between minimal initial configuration energy vs. minimal search energy must be further explored in light of particular applications.

### Conclusions

We have begun to explore how space-filling curves can enhance the efficiency and robustness of geographic search by robot collectives. Initial results have been quite promising, especially in applications such as mine-clearing where exhaustive search is necessary. Much work still needs to be done, but it appears that combining this kind of search with more information-exploitive search algorithms [GR98b] may be extremely useful in general real-world search situations.

We hope to have movies of the robot simulations, as well as an expanded version of this report, online soon at <http://www.sandia.gov/aisl/robotics/>.

### References

- [B69] Bially, T. Space-Filling Curves: Their Generation and Their Application to Bandwidth Reduction. *IEEE Transactions on Information Theory*, V 15 No. 6, November 1969, pp. 658-664.
- [CFKM95] Cao, U., Fukunaga, A., Kahng, A., and Meng, F. 1995. Cooperative mobile robotics: Antecedents and directions. *Proc. of IEEE/RSJ IROS*, pp. 226-234.
- [CP97] Choset, H., and Pignon, P., Coverage Path Planning: The Boustrophedon Cellular Decomposition. International Conference on Field and Service Robotics, Canberra, Australia, 1997. Also available at <http://voronoi.sbp.ri.cmu.edu/~choset>
- [G93] Gage, D., How to communicate to zillions of robots. *Mobile Robots VIII, SPIE*, 250-257, 1993.
- [G97] Gilbert, W., *A Cube-Filling Hilbert Curve*, <http://math.uwaterloo.ca/~wgilbert/Research/HilbertCurve/HilbertCurve.html>
- [GR98a] Goldsmith, S., and Robinett, R. *Collaborative Search by Mobile Robots, Part I: Problem Definition*.

<sup>8</sup> It can know where the next robot's starting point is either by memorizing the location during the initial configuration communication, or by having each robot drop a marker [R97] on the ground at its own starting point. In the latter case, initial communication may be unnecessary.

- Sandia National Laboratories Technical Report. Preprint at <http://www.sandia.gov/aisl/robotics/collaboration>.
- [GR98b] Goldsmith, S., and Robinett, R., *Collective Search by Mobile Robots Using Alpha-Beta Coordination*. Collective Robotics Workshop '98, Agent World, Paris, 1998.
- [HTL96] Hert, S., Tiwari, S., and Lumelsky, V. A Terrain-Covering Algorithm for an AUV. *Autonomous Robots*, 3:91-119, 1996.
- [K95] Kesteloot, L. *Fault-Tolerant Distributed Consensus*. <http://tofu.alt.net/~lk/290.paper/290.paper.html>
- [KOAY95] Kurabayashi, D., Ota, J., Arai, T. and Yoshida, E. Cooperative Sweeping by Multiple Mobile Robots, *1996 IEEE Intl. Conf. on Robotics and Automation*, pp. 1744-1749, 1996.
- [LSP82] Lamport, L., Shostak, R., and Pease, M. The Byzantine Generals Problem, *ACM Transactions on Programming Languages and Systems* 4, 3 (July 1982), 382--401.
- [MJFS96] Moon, B., Jagadish, H. V., Faloutsos, C. and Saltz, J. *Analysis of the Clustering Properties of Hilbert Space-filling Curve*, 1996 University of Maryland Technical Report CS-TR-3611, <http://www.cs.umd.edu/TR/UMCP-CSD:CS-TR-3611>
- [McW97] McWhorter, W., *Fractint L-System Variations*, <http://spanky.triumf.ca/www/fractint/lsys/variations.html>
- [PLF91] Prusinkiewicz, P., Lindenmayer, A. and Fracchia, F. D., Synthesis of Space-filling Curves on the Square Grid. *Fractals in the Fundamental and Applied Sciences*, edited by Peitgen, H.-O. et al., Elsevier Science Publishers, 1991.
- [R97] Russell, R. A., Heat Trails as Short-Lived Navigational Markers for Mobile Robots, *1997 IEEE Intl. Conf. on Robotics and Automation 1997*, pp. 3534-3539, 1997.
- [SLP83] Stevens, R. J., Lehar, A. F. and Perston, F. H. Manipulation and presentation of multidimensional image data using the Peano scan. *IEEE Trans. on Pattern Analysis and Machine Intelligence*, PAMI-5(5) (1983) pp. 520-526.
- [S94] Sagan, H. *Space-filling Curves*, Springer-Verlag, New York 1994.
- [W97a] Weisstein, E., *Random Walk: 2-D*, <http://www.astro.virginia.edu/~eww6n/math/RandomWalk:2-D.html>
- [W97b] Weisstein, E., *Hilbert Curve*, <http://www.astro.virginia.edu/~eww6n/math/HilbertCurve.html>